


SOLE INVENTOR

Docket No. 20002/16136X2  
P16136X2

"EXPRESS MAIL" mailing label No.  
EV 266323687 US

Date of Deposit: **June 27, 2003**

I hereby certify that this paper (or fee) is  
being deposited with the United States  
Postal Service "EXPRESS MAIL POST  
OFFICE TO ADDRESSEE" service  
under 37 CFR §1.10 on the date  
indicated above and is addressed to:  
Mail Stop Patent Application,  
Commissioner for Patents, P.O. Box  
1450, Alexandria, VA 22313-1450



---

Magda Greer

## APPLICATION FOR UNITED STATES LETTERS PATENT

# SPECIFICATION

TO ALL WHOM IT MAY CONCERN:

Be it known that I, **Mingqiu SUN**, a citizen of United States of  
America, residing at 15360 SW Kiwanda Lane, Beaverton, Oregon, 97007  
have invented new and useful **METHODS AND APPARATUS TO  
PREFETCH MEMORY OBJECTS**, of which the following is a  
specification.

# METHODS AND APPARATUS TO PREFETCH MEMORY OBJECTS

## RELATED APPLICATION

**[0001]** This patent issued from a continuation-in-part of U.S. Application Serial No. 10/424,356, which was filed on April 28, 2003.

## FIELD OF THE DISCLOSURE

**[0002]** This disclosure relates generally to memory management, and, more particularly, to methods and apparatus to prefetch memory objects.

## BACKGROUND

**[0003]** Programs executed by computers and other processor based devices typically exhibit repetitive patterns. It has long been known that identifying such repetitive patterns provides an opportunity to optimize program execution. For example, software and firmware programmers have long taken advantage of small scale repetitive patterns through the use of iterative loops, etc. to reduce code size, control memory allocation and perform other tasks seeking to optimize and streamline program execution.

**[0004]** Recently, there has been increased interest in seeking to identify larger scale repetition patterns in complicated workloads such as, for example, managed run-time environments and other server-based applications, as a mechanism to optimize handling of those workloads. For instance, it is known that a workload may be conceptualized as a series of macroscopic transactions. As used herein, the terms macroscopic transaction and sub-

transaction refer to a business level transaction and/or an application software level transaction. For instance, the workload of a server at an Internet retailer such as Amazon.com may be conceptualized as an on-going sequence of macroscopic transactions and sub-transactions such as product display, order entry, order processing, customer registration, payment processing, etc. Moving to a more microscopic level, each of the macroscopic transactions in the workload may be seen as a series of program states. It is desirable to optimize the execution of workloads by, for example, reducing the time it takes the hosting computer to transition between and/or execute macroscopic transactions and/or program states. Therefore, there is an interest in identifying repetition patterns of program states in macroscopic transactions in the hope of predicting program state transitions, optimizing the execution of macroscopic transactions and/or program states, and increasing the throughput of the workload associated with such transactions.

**[0005]** There have been attempts to exploit repetitive structures such as loops to, for example, prefetch data to a cache. However, those prior art methodologies have been largely limited to highly regular and simple workloads such as execution of scientific codes. Effectively predicting program states and/or macroscopic transactions for larger, more complicated workloads remains an open problem.

## BRIEF DESCRIPTION OF THE DRAWINGS

**[0006]** FIG. 1 is a schematic illustration of an example apparatus to detect patterns in programs.

**[0007]** FIG. 2 is a more detailed schematic illustration of the example state identifier of FIG. 1.

**[0008]** FIG. 3 is a schematic illustration of an example trace.

**[0009]** FIG. 4 is a diagram illustrating an example manner in which the signature developer and the weight assigning engine of FIG. 3 may operate to develop signatures.

**[0010]** FIG. 5 illustrates an example data structure which may be created for each identified state in the program.

**[0011]** FIG. 6 is a more detailed schematic illustration of the example predictor of FIG. 1.

**[0012]** FIG. 7 is a chart graphing example entropy values calculated by the entropy calculator of FIG. 6.

**[0013]** FIG. 8 is a flow chart illustrating example machine readable instructions for implementing the trace sampler of the apparatus of FIG. 1.

**[0014]** FIGS. 9A-9C are flowcharts illustrating example machine readable instructions for implementing the state identifier and predictor of the apparatus of FIG. 1.

**[0015]** FIG. 10 is a schematic illustration of an example apparatus to prefetch memory objects.

**[0016]** FIG. 11 illustrates an example data structure which may be created by the apparatus of FIG. 10 for each identified state in the program.

**[0017]** FIG. 12 is a schematic illustration of an example program execution path.

**[0018]** FIGS. 13A-13C are flowcharts illustrating example machine readable instructions for implementing the state identifier, the memory state monitor, the prefetcher, and the predictor of the apparatus of FIG. 10.

**[0019]** FIG. 14 is a flowchart illustrating example machine readable instructions which may be executed to implement a first example prefetching strategy.

**[0020]** FIG. 15 is a flowchart illustrating example machine readable instructions which may be executed to implement a second example prefetching strategy.

**[0021]** FIG. 16 is a schematic illustration of an example computer which may execute the programs of FIGS. 8 and 9A-9C to implement the apparatus of FIG. 1, and/or which may execute the programs of FIG. 8, FIGS. 13A-13C, FIG. 14 and/or FIG. 15 to implement the apparatus of FIG. 10.

## DETAILED DESCRIPTION

**[0022]** As mentioned above, real world server applications typically exhibit repetitive behaviors. These repetitive behaviors are usually driven by local or remote clients requesting performance of tasks or business transactions defined by the application program interface (API) of the host site. Since the range of tasks available to the clients is limited, the client calls into the API manifest themselves as repetitive program execution patterns on the hosting server. As explained below, this type of repetitiveness provides efficiency opportunities which may be exploited through microprocessor architecture and/or software.

**[0023]** The basic unit of repetition within these repetitive program execution patterns is a macroscopic transaction or sub-transaction. A macroscopic transaction or sub-transaction may be thought of as one pathlength measured by instructions. The pathlength of such a transaction or sub-transaction is typically, for example, in the range of  $10^4$  to  $10^6$  instructions.

**[0024]** Each transaction or sub-transaction includes one or more program states. A program state is defined as a collection of information (e.g., a series of memory addresses and/or a series of instruction addresses) occurring in a given time window. A program state may be a measurement-dependent and tunable property. On the other hand, a transaction or a sub-transaction is typically an intrinsic property of a workload.

**[0025]** FIG. 1 is a schematic illustration of an example apparatus 10 to predict program states of an executing program and/or to identify macroscopic transactions of the program. For the purpose of developing a trace of a program of interest, the apparatus 10 is provided with a trace sampler 12. The trace sampler 12 operates in a conventional fashion to develop any type of trace of the program of interest. For example, the trace sampler 12 may employ a hardware counter such as a processor counter and/or software instrumentation such as managed run-time environment (MRTE) instrumentation to gather trace data from the executing program. For instance, the trace sampler 12 may capture the instruction addresses appearing in the program counter of a processor to create an instruction address trace. By way of another example, the trace sampler 12 may

snoop an address bus associated with the cache of a processor to create a memory address trace. Persons of ordinary skill in the art will readily appreciate that many other techniques can be used to create the same or different types of traces. For instance, the trace sampler 12 could alternatively be configured to create a basic block trace.

**[0026]** In order to identify a sequence of program states from the trace generated by the trace sampler 12, the apparatus 10 is further provided with a state identifier 14. As will be appreciated by persons of ordinary skill in the art, the state identifier 14 may identify the states within the trace created by (or being created by) the trace sampler 12 in any number of ways. In the illustrated example, the state identifier 14 identifies the program states by comparing adjacent sets of data at least partially indicative of entries appearing in the trace. To make this comparison more manageable, the illustrated state identifier 14 translates the sets into bit vectors which function as short hand proxies for the data in the sets. The illustrated state identifier 14 then compares the bit vectors of adjacent sets and determines if the difference between the bit vectors is sufficient to indicate that a new state has occurred. Each of the sets of data may comprise sequential groups of entries in the trace. Either all of the entries in the trace may be used, or a subset of the entries may be used (e.g., every tenth entry may be used) to create the sets. Further, either a fraction of the entries selected to be in the set (e.g., the last eight bits) or the entire portion of the entry (e.g., all of the bits in the entry) may be used to create the bit vectors. Persons of ordinary skill in the art will readily appreciate that adjusting the resolution of the sets (e.g., by adjusting the

number of entries skipped in creating the sets and/or by adjusting the amount or location of the bits of the entries in the trace that are used to create the bit vectors), may adjust the identities of the program states that are identified by the state identifier 14. Thus, the program state definitions are measurement-dependent and tunable.

**[0027]** An example state identifier 14 is shown in FIG. 2. In the illustrated example, the state identifier 14 includes a signature developer 16 to develop possible state signatures from the sets of entries in the trace. To better illustrate the operation of the signature developer 16, consider the example trace shown in FIG. 3. In the example of FIG. 3, the trace 18 comprises a sequential series of entries representative in some fashion of a characteristic of the computer and/or a component thereof that changes over time as a result of executing the program of interest. For example, the entries may be instruction addresses appearing in the program counter of a processor, memory addresses appearing on an address bus of the cache associated with the processor, or any other recordable characteristic in the computer that changes as a result of executing the program. Persons of ordinary skill in the art will appreciate that the entries may be complete addresses, portions of complete addresses, and/or proxies for complete or partial addresses. In view of the broad range of possibilities for the types of data logged to create the entries of the trace 18, FIG. 3 generically describes these entries by the symbol “A” followed by a number. The number following the symbol “A” serves to uniquely distinguish the entries. To the extent execution of the program of interest causes the monitored characteristic used to create



the trace to have the same value two or more times, the trace 18 will include the same entry two or more times (e.g., entry A5 appears twice in the trace 18). The number following the symbol “A” may indicate a relative position of the entry relative to the other entries. For example, if the trace 18 is an instruction address trace, each number following the letters may represent a location in memory of the corresponding address. For simplicity of explanation, unless otherwise noted, the following example will assume that the trace 18 is an instruction address trace reflecting the full memory addresses of the instructions executed by a processor running a program of interest.

**[0028]** The primary purpose of the signature developer 16 is to create proxies for the entries in the trace 18. In particular, the entries in the trace 18 may contain a significant amount of data. To convert these entries into a more manageable representation of the same, the signature developer 16 groups the entries into sets 26 and converts the sets 26 into possible state signatures 28. In the illustrated example, the possible state signatures 28 are bit vectors. The sets 26 may be converted into bit vectors 28 as shown in FIG. 4.

**[0029]** In the example of FIG. 4 a random hashing function 30 is used to map the entries in a set 26 to an n-bit vector 28. In the example of FIG.4, the value “B” 32 defines the resolution of the model (e.g., the number of entries in the set 26 that are skipped (if any) and/or processed by the hash function 30 to generate the n-bit vector 28). The basic use of a hash function 30 to map a set of entries from a trace 18 into a bit vector is well known to persons of ordinary skill in the art (see, for example, Dhodapkar & Smith, “Managing Multi-Configuration Hardware Via Dynamic Working Set

Analysis,” <http://www.cae.wisc.edu/~dhodapka/isca02.pdf>) and thus, in the interest of brevity, will not be further explained here. The interested reader can refer to any number of sources, including, for example, the Dhodapkar & Smith article mentioned above, for further information on this topic.

**[0030]** For the purpose of weighting the members of the sets 26 such that later members have greater weight than earlier members of the set 26 when mapping the set 26 of entries to the bit vector signature 28, the apparatus 10 is further provided with a weight assigning engine 34. As shown in the example mapping function of FIG. 4, the weight assigning engine 34 applies an exponential decay function 36 (e.g.,  $f_1 = e^{-t/T}$  where  $t$  =time and  $T$  =half lifetime) to the entries in a set 26 prior to operating on the set 26 with the hashing function 30. The exponential decay function 36 is applied to the entries in the set 26 of entries so that, when the hashing function 30 is used to convert the set 26 into a possible state signature 28, the latest entries in the set 26 have a greater impact on the values appearing in the possible state signature 28 than earlier values in the set 26. Persons of ordinary skill in the art will appreciate that, as with other structures and blocks discussed herein, the weight assigning engine 34 is optional. In other words, the exponential decay function 36 shown in FIG. 4 may be optionally eliminated.

**[0031]** As explained above, the illustrated signature developer 16 operates on sequential sets 26 of the entries appearing in the trace 18 to create a series of bit vectors 28 corresponding to those sets 26. Persons of

ordinary skill in the art will readily appreciate that the signature developer 16 may group the entries in the trace 18 into sets 26 in any number of ways. However, in the illustrated example, the signature developer 16 creates the sets 26 such that adjacent sets 26 overlap (i.e., share at least one entry). In other words, the signature developer 16 uses a sliding window to define a series of overlapping sets 26. The number of entries in the trace 18 that are shared by adjacent sets 26 (i.e., the intersection of adjacent sets) may be as small as one element or as large as all but one element (see, for example, the overlapping sets 26 in FIG. 4). In examples in which the signature developer 16 creates adjacent intersecting sets 26, it is particularly advantageous to also use the weight assigning engine 34 such that the possible state signatures 28 created by the signature developer 16 are more responsive to the newer non-overlapping entries than to the overlapping entries and the older non-overlapping entries.

**[0032]** In order to identify program states based on the possible state signatures 28, the apparatus 10 is further provided with a state distinguisher 38. In the illustrated example, the state distinguisher 38 begins identifying program states by selecting one of the possible state signatures 28 as a first state signature 40 (e.g., State 1 in FIG. 3) to provide a reference point for the remainder of the analysis. Typically, the first possible state signature 28 (e.g., PS1 in FIG. 3) in the sequence of possible state signatures (e.g., PS1-PSN) is, by default, selected as the first state signature 40, but persons of ordinary skill in the art will readily appreciate that this selection is arbitrary and another one

of the possible state signatures 28 (e.g., PS2-PSN) may alternatively be used as the first state signature 40.

**[0033]** Once a first state signature 40 is selected, the state distinguisher 38 compares the first state signature 40 to a next subsequent one of the possible state signatures 28 (e.g., PS2). For example, if the first state signature 40 is the first possible state signature, the first state signature 40 may be compared to the second possible state signature PS2 in the list of possible state signatures 28. If the next subsequent state signature 28 (e.g., PS2) differs from the first state signature 40 by at least a predetermined amount, there has been sufficient change in the measured parameter used to create the trace 18 to designate the corresponding program as having entered a new program state. Accordingly, the state distinguisher 38 identifies the subsequent possible state signature 28 (e.g., PS2) as a second state signature.

**[0034]** If, on the other hand, the subsequent state signature 28 (e.g., PS2) does not differ from the first state signature 40 by at least a predetermined amount, there has not been sufficient change in the measured parameter used to create the trace 18 to designate the corresponding program as having entered a new program state. Accordingly, the state distinguisher 38 discards the possible state signature 28 (e.g., PS2), skips to the next possible state signature 28 (e.g., PS3), and repeats the process described above by comparing the first state signature 40 to the next possible state signature 28 (e.g., PS3). The state distinguisher 38 continues this process of sequentially comparing possible state signatures 28 (e.g., PS2-PSN) to the first state signature 40 until a possible state signature 28 (e.g., PS4) is identified that

differs from the first state signature 40 by at least the predetermined amount.

When such a possible state signature (e.g., PS4) is identified, the state distinguisher 38 designates that possible state signature (e.g., PS4) as the second state signature (e.g., State 2). All intervening possible state signatures 28 (e.g., PS2-PS3) are not used again, and, thus, may be discarded.

**[0035]** Once the second state (e.g., State 2) is identified, the state distinguisher 38 then begins the process of comparing the second state signature (e.g., PS4) to subsequent possible state signatures (e.g., PS5, etc.) to identify the third state (e.g., State 3) and so on until all of the possible state signatures (e.g., PS2-PSN) have been examined and, thus, all of the program states (State 1 – State N) occurring during the current execution of the program have been identified. Example program states (i.e., State 2 – State N) appearing after the first program state 40 are shown in FIG. 3. As shown in that example, any number of program states may occur and/or reoccur any number of times depending on the program being analyzed.

**[0036]** Persons of ordinary skill in the art will appreciate that there are many possible ways to compare the state signatures (e.g., State 1 – State N) to subsequent possible state signatures (e.g., PS2 – PSN) to determine if a new program state has been entered. Such persons will further appreciate that there are many different thresholds that may be used as the trigger for determining that a new state has been entered. The threshold chosen is a determining factor in the number and definitions of the states found in the program. In the illustrated example, the threshold difference required between signatures to declare a new program state is the Hamming distance. Thus, if the difference

between a state signature (e.g., State 1) and a possible state signature (e.g., PS2) satisfies the following equation, then a new program state has been entered:

$$\text{[0037] } \Delta = |\text{State Signature XOR Possible State Signature}| / |\text{State Signature OR Possible State Signature}|$$

**[0038]** In other words, a new state has been entered in the example implementation if the set of bit values appearing in only one of: (a) the current state signature and (b) a possible state signature (i.e., the set of differences) divided by the set of all members appearing in either (a) the current state signature and/or (b) the possible state signature (i.e., the total set of members (e.g., logic one values appearing in the bit vectors)) is greater than a predetermined value (e.g.,  $\Delta$ ).

**[0039]** To manage data associated with the states identified by the state distinguisher 38, the apparatus 10 is further provided with a memory 44 (see FIG. 1). The memory 44 of the illustrated example is configured as a state array including a plurality of state data structures, wherein each data structure corresponds to a unique program state. As will be appreciated by persons of ordinary skill in the art, the state data structures and the state array 44 may be configured in any number of manners. In the illustrated example, the state array 44 is large enough to contain four hundred state data structures and each data structure in the state array includes the following fields: (a) the state signature of the corresponding program state, (b) an age of the corresponding program state, (c) a usage frequency of the corresponding program state, (d) an entropy value of the corresponding state, and (e) a sub-array containing a

set of probabilities of transitioning from the corresponding program state to a set of program states.

**[0040]** An example state data structure is shown in FIG. 5. The state signature field may be used to store the bit vector signature (e.g., State 1 – State N) of the state corresponding to the data structure. The age field may be used to store a value indicative of the time at which the corresponding state was last entered. Because the state array is finite, the age field may be used as a vehicle to identify stale state data structures that may be over written to store data for a more recently occurring state data structure. The usage frequency field may be used to store data identifying the number of times the corresponding state has been entered during the lifetime of the data structure. The entropy value field may be used to store data that may be used to identify the end of a macroscopic transaction. The set of probabilities sub-array may be used to store data indicating the percentage of times program execution has entered program states from the program state corresponding to the state data structure during the lifetime of the state data structure. For example, each data structure may store up to sixteen sets of three fields containing data indicating a name of a program state to which the program state corresponding to the state data structure has transitioned in the past, the relative time(s) at which those transitions have occurred, and the percentage of times that the program state corresponding to the state data structure has transitioned to the state identified in the first field of the set of fields.

**[0041]** In order to determine entropy values associated with the program states identified by the state identifier, the apparatus 10 is further

provided with a predictor 46. As explained below, in the illustrated example, the predictor 46 uses the entropy values to identify an end of a macroscopic transaction.

**[0042]** An example predictor 46 is shown in greater detail in FIG. 6. To calculate probabilities of transitioning from one of the program states to another of the program states, the predictor 46 is provided with a state transition monitor 48. Whenever a program state transition occurs (i.e., whenever the state of the program changes from one state to another), the state transition monitor 48 records the event in the sub-array of the state data structure corresponding to the program state that is being exited. In particular, the state transition monitor 48 records data indicating the name of the array transitioned to and the time (or a proxy for the time) at which the transition occurred. The time (or a proxy for the time) at which the transition occurred is recorded because, in the illustrated example, the state transition monitor 48 calculates the probabilities as exponential moving averages. Thus, instead of merely averaging the entries in the sub-array of the state data structure to calculate the probabilities of transitioning between specific states based on past performance, the state transition monitor 48 weights the entries in the sub-array of the state data structure based on their relative times of occurrence by multiplying those entries by an exponential function. As a result of this approach, entries in the sub-array which occur later in time have greater weight on the probability calculations than entries which occur earlier in time, and the state transition monitor 48 can, thus, identify changing patterns in the probabilities more quickly than an approach using straight moving averages.



**[0043]** To convert the probabilities calculated by the state transition monitor 48 into entropy values, the apparatus 10 is further provided with an entropy calculator 50. The entropy value of a given state is the transitional uncertainty associated with that state. In other words, given the past history of a current state, the entropy value quantifies the informational uncertainty as to which program state will occur when the current program state ends. For instance, for a given program state that has a past history of transitioning to a second program state and a third program state, the entropy calculator 50 converts the probabilities to entropy values for the given program state by calculating a sum of (1) a product of (a) a probability of transitioning from the subject program state to the second program state and (b) a logarithm of the probability of transitioning from the subject program state to the second program state, and (2) a product of (a) a probability of transitioning from the subject program state to the third program state and (b) a logarithm of the probability of transitioning from the subject program state to the third program state. Stated another way, for each state data structure in the state array 44, the entropy converter 50 calculates an entropy value in accordance with the well known Shannon formula:

$$\mathbf{[0044] \quad H = -K \sum (P_i * \log P_i),}$$

**[0045]** where H is the entropy value, K is a constant and  $P_i$  is the probability of transitioning from the current state (i.e., the state associated with the state data structure) to state “i” (i.e., the states identified in the sub-array of the data structure of the current state). The entropy value of each state

identified in the executing program is stored in the data structure of the corresponding state (see FIG. 5).

**[0046]** In order to predict the next probable program state to be transitioned to from the current state, the predictor 46 further includes an event predictor 54. The event predictor 54 compares the probabilities appearing in the sub-array of the data structure of the current program state to determine the next most probable state or states. The next most probable state(s) are the state(s) that have the highest probability values.

**[0047]** The event predictor 54 also functions to identify a macroscopic transaction based on the entropy value associated with the current program state. Viewed from a macroscopic application logic level, one can observe a link to the calculated entropy value ( $H$ ), which is a microscopic trace property. When a new business transaction starts, program execution typically follows a relatively well-defined trajectory with low entropy. However, as program execution reaches the last program state in a macroscopic transaction, the entropy value spikes as there is maximum uncertainty about the possible next program state to which the program will transition. In other words, within a macroscopic transaction, there are typically repetitive sequences of program states. By observing past behavior between program states, one can detect these patterns and use them to predict future behavior. In contrast, the order of macroscopic transactions has a higher degree of randomness than the order of program states within a macroscopic transaction because the order in which macroscopic transactions are executed depends on the order in which requests for transactions are received from third parties and is, thus, substantially

random. To make this point clearer, consider an on-line retailer. The server of the on-line retailer receives requests from a number of different customers and serializes those requests in a generally random fashion in a queue. The order in which the requests are handled is, thus, random. However, once the server begins serving a request, it will generally process the entire transaction before serving another transaction from the queue. As a result, the program state at the end of a macroscopic transaction typically has a high entropy value (i.e., there is a high level of uncertainty as to which program state will be entered), because there is a high level of uncertainty as to which macroscopic transaction will follow the current macroscopic transaction that just completed execution. Consequently, the last program state in a macroscopic transaction is characterized by a spike in its entropy value relative to the surrounding entropy values. In other words, the entropy value of the last program state of a macroscopic transaction is typically a relative maximum as compared to the entropy values of the program states immediately preceding and following the last program state.

**[0048]** The event predictor 54 takes advantage of this characteristic by using this entropy spike as a demarcation mark for the end of a macroscopic transaction. A macroscopic transaction may thus be defined as an ordered sequence of program states with an entropy-spiking ending state. A macroscopic transaction maps to a business or application software transaction, which is an intrinsic property of a workload. The same macroscopic transaction may contain different sets of program states, which are measurement-dependent properties of a workload that can be tuned

through the transition threshold value. One caution, however, is that repeatable sub-transactions that may not be significant to high level business logic may also end at a program state exhibiting a spiking entropy value and, thus, may be mis-identified as a macroscopic transaction. This mis-identification is not a problem in practical cases such as performance tuning of a program because sub-transactions with large transitional uncertainty behave like transactions for all practical purposes.

**[0049]** As stated above, the event predictor 54 identifies a spike in the entropy values of a series of program states as an end of a macroscopic transaction. Persons of ordinary skill in the art will appreciate that the event predictor 54 may use any number of techniques to identify a spike in the entropy values. For example, the event predictor 54 may compare the entropy value of the current state to the entropy value of the previous state and the entropy value of the following state. If the entropy value of the current state exceeds the entropy value of the previous state and the entropy value of the following state, the entropy value of the current state is a relative maximum (i.e., a spike) and the current state is identified as the end of a macroscopic transaction. Otherwise, it is not a relative maximum and the current state is not identified as the end of a macroscopic transaction.

**[0050]** A chart illustrating a graph of example entropy values calculated by the entropy calculator 50 is shown in FIG. 7. In the chart of FIG. 7, instead of using the signatures to index the program states, we use the first discovery time of each program state as its unique index. These first discovery times are used as the ordinates of the Y-axis in FIG. 7. (The

ordinates of the Y-axis also represent entropy values as explained below.)

Memory accesses are used as the ordinates of the X-axis of FIG. 7. The memory accesses are a proxy to time.

**[0051]** The chart of FIG. 7 includes two graphs. One of the graphs represents the program states that are entered over the time period at issue. The other graph represents the entropy values of the corresponding program states over that same time period. As can be seen by examining FIG. 7, each state in the graph (i.e., each data point represented by a diamond ♦) is positioned in vertical alignment with its corresponding entropy value (i.e., each data point represented by a square ■). As can also be seen in FIG. 7, the entropy values spike periodically. Each of these spikes in the entropy values represents an end of a macroscopic transaction.

**[0052]** Flowcharts representative of example machine readable instructions for implementing the apparatus 10 of FIG. 1 are shown in FIGS. 8 and 9A-9C. In this example, the machine readable instructions comprise a program for execution by a processor such as the processor 1012 shown in the example computer 1000 discussed below in connection with FIG. 10. The program may be embodied in software stored on a tangible medium such as a CD-ROM, a floppy disk, a hard drive, a digital versatile disk (DVD), or a memory associated with the processor 1012, but persons of ordinary skill in the art will readily appreciate that the entire program and/or parts thereof could alternatively be executed by a device other than the processor 1012 and/or embodied in firmware or dedicated hardware in a well known manner. For example, any or all of the trace sampler 12, the state identifier 14, the

predictor 46, the weight assigning engine 34, the signature developer 16, the state distinguisher 38, the state transition monitor 48, the entropy calculator 50, and/or the event predictor 54 could be implemented by software, hardware, and/or firmware. Further, although the example program is described with reference to the flowcharts illustrated in FIGS. 8 and 9A-9C, persons of ordinary skill in the art will readily appreciate that many other methods of implementing the example apparatus 10 may alternatively be used. For example, the order of execution of the blocks may be changed, and/or some of the blocks described may be changed, eliminated, or combined.

**[0053]** The program of FIG. 8 begins at block 100 where the target program begins execution. While the target program executes, the trace sampler 12 creates one or more traces 18 of one or more properties of the executing program (block 102). For example, the trace sampler 12 may generate an instruction address trace, a memory address trace, a basic block trace, and/or any other type of trace. Control proceeds from block 102 to block 104.

**[0054]** If a trace processing thread has already been invoked (block 104), control proceeds from block 104 to block 106. If the trace 18 of the program is complete (block 106), the program of FIG. 8 terminates. Otherwise, if the trace 18 of the program is not complete (block 106), control returns to block 102 where the recording of the trace 18 continues.

**[0055]** If the trace processing thread has not already been invoked (block 104), control proceeds to block 108. At block 108, the trace processing thread is initiated. Control then returns to block 106. As explained above, the

program will terminate if the trace 18 is complete (block 106) or continue to generate the trace 18 (block 102) if the trace 18 has not yet been completed. Thus, control continues to loop through blocks 100-108 until the target program stops executing and the trace 18 is complete.

**[0056]** An example trace processing thread is shown in FIGS. 9A-9C. The illustrated trace processing thread begins at block 120 where the signature developer 16 obtains a set 26 of entries from the trace 18 created by the trace sampler 12. As explained above, the sets 26 of entries may be created in any number of ways to include any number of members. In the example of FIG. 3, each of the sets 26 include a series of sequential entries (i.e., no entries are skipped), and adjacent sets overlap (i.e., at least one of the entries is used in two adjacent sets. However, sets which skip some entries in the trace 18 and/or which do not overlap could alternatively be employed.

**[0057]** Once the entries to create a set 26 are retrieved from the trace 18 (block 120), the weight assigning engine 39 adjusts the values of the retrieved entries such that later entries are given greater weight than earlier entries (block 122). For example, the weight assigning engine 34 may apply an exponential decay function 36 (e.g.,  $f_1 = e^{-t/T}$ ) to the entries in the set (block 122).

**[0058]** Once the values of the entries have been weighted by the weight assigning engine 34 (block 122), the signature developer 16 maps the entries in the set 26 to an n-bit vector to create a possible state signature 28 for the set 26 (block 124). As explained above, the mapping

of the entries in the set 26 to the possible state signature 28 may be performed using a hashing function.

**[0059]** After the possible state signature 28 is generated (block 124), the state distinguisher 38 determines whether the possible state signature 28 is the first possible state signature (block 126). If it is the first possible state signature (block 126), the first possible state signature is, by default, defined to be the first state signature. Thus, the state distinguisher 38 sets a current state signature variable equal to the possible state signature 28 (block 128) and creates a state data structure in the state array 44 for the first state (block 130). An example state data structure is shown in FIG. 5. The state distinguisher 38 may create the state data structure by creating the fields shown in FIG. 5, by writing the current state signature into the state signature field of the new state data structure, by setting the age field of the new state data structure equal to the current time or a proxy for the current time, and by setting the entropy field and the probability sub-array fields equal to zero.

**[0060]** The signature developer 16 then collects the next set 26 of entries for creation of a possible state signature 28 (block 132). In the illustrated example, the sets 26 used by the signature developer 16 to create the possible signatures 28 are overlapping. Thus, the signature developer 16 may create the next set 26 of entries by dropping the oldest entr(ies) from the last set 26 of entries and adding a like number of new entr(ies) to create a new current set 26 (block 132). Control then returns to block 122 where the entries in the new current set are weighted as explained above.



**[0061]** When at block 126, the current possible state signature is not the first possible state signature, control will skip from block 126 to block 134 (FIG. 9B). At block 134, the state distinguisher 38 calculates the difference between the current state signature (i.e., the value in the current state signature variable mentioned above), and the current possible state signature. The state distinguisher 38 then compares the computed difference to a threshold (e.g., the Hamming difference). If the computed difference exceeds the threshold (block 136), a program state change has occurred and control proceeds to block 138. If the computed difference does not exceed the threshold (block 136), the signature developer 16 collects the next set 26 of entries for creation of a possible state signature 28 (block 132, FIG. 9A) and control returns to block 122 as explained above. Thus, control continues to loop through blocks 122-136 until a program state change occurs.

**[0062]** Assuming for purposes of discussion that a program state change has occurred (block 136), the state distinguisher 38 sets the current state signature variable equal to the current possible state signature 28 (block 138). The state distinguisher 38 then examines the signatures present in the state array 44 to determine if the current state signature corresponds to the signature of a known state (block 140). If the current state signature is a known state signature, control advances to block 160 (FIG. 9C). Otherwise, if the current state signature is not a known state signature (i.e., the current state signature does not correspond to a state

already existing in the state array 44), control advances to block 142 (FIG. 9B).

**[0063]** Assuming for purposes of discussion that the current state signature is not a known state signature (e.g., the current program state is a new program state) (block 140), the state distinguisher 38 creates a state data structure in the state array 44 for the first state (block 142) as explained above in connection with block 130.

**[0064]** The state transition monitor 48 then updates the last state's probability sub-array to reflect the transition from the last state to the new current state (block 144). Control then proceeds to block 146 where the state distinguisher 38 determines if the state array 44 has become full (i.e., if the newly added data structure used the last available spot in the state array). If the state array 44 is not full, control returns to block 132 (FIG. 9A) where the signature developer 16 collects the next set 26 of entries for creation of a possible state signature 28. Control then returns to block 122 as explained above.

**[0065]** If the state array is full (block 146), control advances to block 150 (FIG. 9B) where the state distinguisher 38 deletes the stalest state data structure from the state array 44. The stalest state data structure may be identified by comparing the usage fields of the state data structures appearing in the state array 44. Once the stalest state data structure is eliminated (block 150), control advances to block 132 where the signature developer 16 collects the next set 26 of entries for creation of a

possible state signature 28. Control then returns to block 122 as explained above.

**[0066]** Assuming that the current state signature is a known state signature (block 140), control proceeds to block 160 (FIG. 9C). The state transition monitor 48 then updates the last state's probability sub-array to reflect the transition from the last state to the new current state (block 160). Control then proceeds to block 162 where the entropy calculator 50 calculates the entropy value of the current state. As explained above, the entropy value may be calculated in many different ways. For instance, in the illustrated example, the entropy value is calculated using the Shannon formula.

**[0067]** Once the entropy value is calculated (block 162), the event predictor 54 identifies the next most probable state(s) (block 164) by, for example, comparing the values in the probability sub-array of the state data structure of the current state. The event predictor 54 may then examine the entropy values of the last few states to determine if an entropy spike has occurred (block 168). If an entropy spike is identified (block 168), the event predictor 54 identifies the program state corresponding to the entropy spike as the last state of a macroscopic transaction (block 170). If an entropy spike is not identified (block 168), the end of a macroscopic transaction has not occurred. Accordingly, control skips block 170 and returns to block 132 (FIG. 9A).

**[0068]** Irrespective of whether control reaches block 132 via block 170 or directly from block 168, at block 132 the signature developer 16

collects the next set 26 of entries for creation of a possible state signature 28. Control then returns to block 122 as explained above. Control continues to loop through blocks 122-170 until the entire trace 18 has been processed. Once the entire trace 18 has been processed, the trace processing thread of FIGS. 9A-9C terminates.

**[0069]** Persons of ordinary skill in the art will readily appreciate that the above described program state identification framework may be employed (in some cases, with modifications) to achieve various performance enhancements. For example, the above described framework may be modified to detect program state execution patterns and to leverage those patterns to achieve more efficient memory usage. To further elucidate this point, an example apparatus 300 to prefetch memory objects to reduce cache misses is shown in FIG. 10.

**[0070]** The example apparatus 300 of FIG. 10 utilizes some of the same structures as the apparatus 10 of FIG. 1. Indeed, the illustrated apparatus 300 incorporates all of the structures of the apparatus of FIG. 1 (as shown in FIG. 10 by the structures bearing the same names and/or reference numbers as the corresponding structures in FIG. 1), and adds other structures to perform additional functionality. However, persons of ordinary skill in the art will appreciate that, if desired, structures appearing in the example apparatus 10 may be eliminated from the example apparatus 300 of FIG. 10.

**[0071]** Since there is overlap between the structures and functionality of the example apparatus 10 and the example apparatus 300, in the interest of brevity, descriptions of the overlapping structures and functions will not be

fully repeated here. Instead, the interested reader is referred to the corresponding description of the example apparatus 10 of FIG. 1 for a complete description of the similar structures appearing in the example apparatus 300 of FIG. 10. To facilitate this process, like structures are labeled with the same names and/or reference numerals in the figures and descriptions of the apparatus 10 and the apparatus 300.

**[0072]** Like the example apparatus 10, the example apparatus 300 includes a trace sampler 12 to develop a trace of a program of interest, and a program state identifier 14 to identify program states from the trace. It also includes a memory/state array 44 to store data structures containing data representative of the states identified by the program state identifier 14. The illustrated apparatus 300 also includes a predictor 46 to predict the next program state(s) that will likely be entered by the executing program and/or to identify the ends of macroscopic transactions.

**[0073]** In the illustrated example, rather than using instruction addresses to create an instruction trace, the trace sampler 12 of the apparatus 300 records the memory addresses (or proxies for the memory addresses) that are issued to retrieve data and/or instructions from the main memory and/or a mass storage device to the cache to create a main memory address trace. Thus, the program states identified by the state identifier 14 of the example apparatus 300 are based on a memory address trace and, consequently, are reflective of patterns in memory accesses, as opposed to patterns in instruction execution as would be the case if the program states were created based on an instruction address trace. Persons of ordinary skill in the art will appreciate,

however, that other types of traces may alternatively be employed to create the trace. For example, an instruction trace may alternatively be used.

**[0074]** Irrespective of the type of trace created by the trace sampler 12, the state identifier 14 analyzes the trace to identify a series of program states as explained above in connection with the example apparatus 10 of FIG. 1. As in the example apparatus 10 of FIG. 1, the program states identified by the state identifier 14 are represented by state data structures stored in the memory/state array 44. As shown in the example of FIG. 11, the state data structures stored in the state array 44 may include the fields described above in connection with the example state data structure shown in FIG. 5 (e.g., state signature, age, usage frequency, entropy, etc.). However, to make it possible to pre-fetch memory objects, the state data structures of FIG. 11 also include one or more fields to store memory profiles for the states identified by the state identifier 14. For example, the data structure discussed above in connection with FIG. 5 may be modified as shown in FIG. 11 to include the memory object references (or proxies for the memory object references which may be reconstructed to form the memory object references) employed to retrieve the memory objects associated with the corresponding program state. As noted above, the memory object references may be memory addresses. Thus, the memory references appended to the example data structure of FIG. 11 may comprise the portion of the memory address trace (or a reference (e.g., a link) to the portion of the memory address trace) corresponding to the subject program state.

**[0075]** As used herein, the term “memory reference” refers to an address (or a proxy for an address) used to retrieve a memory object from a main memory and/or a mass storage device (e.g., a compact disk, a digital versatile disk, a hard disk drive, a flash memory, etc) to the cache memory, and/or an address (or a proxy for an address) used to retrieve a memory object from a mass storage device to the main memory and/or the cache memory. As used herein, the term “memory object” refers to an instruction, part of an instruction, and/or data that is stored in at least one of a main memory, a cache memory, and a mass storage medium. Fetching or prefetching a memory object may involve retrieving a copy of the object from the main memory or mass storage medium and storing one or more copies of the object in one or more levels of the cache, and/or initializing one or more locations in the cache and/or main memory for storage of data and/or instructions.

**[0076]** To associate memory profiles with respective ones of the program states, the example apparatus 300 is further provided with a memory state monitor 302. The memory state monitor 302 populates the state data structures with the memory references (or proxies for the memory references) associated with the corresponding states. Because the memory references of a given state may change to some degree over the lifetime of an executing program, the memory state monitor 302 may be constructed to update the memory profiles as the program being monitored is executed. For instance, the memory state monitor 302 may be structured to filter the memory profiles by adding, deleting, and/or changing one or more of the memory references in

the memory profiles to reflect the memory references most recently associated with the program states. Thus, for example, the memory state monitor 302 may be adapted to filter the memory references included in the memory profiles based on a usage filter model (e.g., the most recently used memory references are kept, while older references are discarded), or based on a miss filter model (e.g., the memory references associated with a cache miss are kept, while references associated with a cache hit are discarded). Usage filtering and/or miss filtering have the advantage of reducing the size of the stored memory profiles. For example, during testing cache miss filtering reduced the required memory data structures by half while achieving substantially the same level of cache performance benefit.

**[0077]** In order to retrieve memory objects that are expected to be used in the near future from a main memory 306 and/or a mass storage device to a cache, the apparatus 300 is further provided with a prefetcher 304. The prefetcher 304 may use any number of strategies as to which memory references should be prefetched at a particular time. For example, the prefetcher 304 may be structured to retrieve the memory references associated with the next most probable state, all of the next probable states, or a subset of the next probable states. The next most probable state(s) are identified by the predictor 46 by reviewing the probabilities appearing in the sub-array of the data structure of the current program state as explained above in connection with the apparatus 10. The prefetcher 304 may identify the memory references required to prefetch the memory objects associated with the next most probable state(s) identified by the predictor 46 from the memory



profile(s) stored in the state data structure(s) of the next most probable state(s) by the memory state monitor 302.

**[0078]** The prefetcher 304 may always retrieve the memory references of the next most probable state or a plurality of the next most probable state(s). Alternatively, the prefetcher 304 may employ the entropy values developed by the predictor 46 to determine the amount of prefetching to be performed. For example, if the entropy value of the current state exceeds a predetermined threshold, there may be so much uncertainty as to the next state that prefetching memory objects may be more likely to pollute the cache than to expedite execution. Accordingly, if the entropy value is sufficiently high, the prefetcher 304 may be adapted to not prefetch any memory objects.

**[0079]** Additionally or alternatively, the prefetcher 304 may be structured to prefetch a different amount of memory objects for different levels of entropy values. For example, if the entropy value of a current program state is less than a predetermined threshold, the prefetcher 304 may prefetch the memory objects associated with a next most probable state. If, on the other hand, the entropy value of the current program state is greater than the same or a different predetermined threshold, the prefetcher 304 may prefetch the memory objects associated with a plurality of next most probable states.

**[0080]** Irrespective of the prefetching strategy employed, it is important to properly time the occurrence of the prefetching operation. If the prefetching operation is performed too early, the prefetched content may be polluted (e.g., destroyed) before they are used and, thus, the prefetched objects

may be unavailable when they are needed. On the other hand, if the prefetching operation is performed too late, the prefetched memory objects may not have reached the cache by the time they are needed. To address this timing concern, the illustrated apparatus 300 performs the prefetching operation near the beginning of a current program state, and the programs states are defined to have a duration that exceeds the latency of the memory 306, but is not long enough to allow the prefetched objects to be polluted before they are needed. To define the program states to have durations meeting these criteria, it may be necessary to tune the threshold difference required between signatures for the state identifier 14 to declare a new program state. Typically, the program states have a duration of a few thousand instructions, which provides sufficient time to prefetch memory objects without causing cache pollution.

**[0081]** An example program execution path illustrating the operation of the example apparatus 300 is shown in FIG. 12. The example of FIG. 12 begins at the start of a macroscopic transaction. At the beginning of that transaction, program state 1 was entered out of a set of probable program states (i.e., states 1, 6, 11, 18 and 21). Upon reviewing the sub-array of the data structure of the current program state, the predictor 46 determines that the next possible states are states 2 and 3, and that state 3 is the next most probable state. (The fact that there are two possible next states illustrates intra-transactional variance.) In the illustrated example, the prefetcher 304 is structured to prefetch the memory references associated with the next most probable state unless the entropy value associated with the current state

exceeds a predetermined value. If the entropy value exceeds that threshold, the prefetcher 304 does not prefetch any memory objects.

**[0082]** In the example of FIG. 12, the entropy value of the current state (i.e., state 1) is sufficiently low to enable prefetching. Accordingly, the prefetcher 304 retrieves the memory profile associated with state 3 from the corresponding state data structure and retrieves the memory objects addressed by the retrieved memory references. The predictor 46 then accesses the sub-array of the data structure associated with state 3 and determines that, based on past performance, the next state (i.e., state 4) is 100% deterministic. Thus, prefetching will be highly effective and the prefetcher 304 accesses the memory profile of state 4 and uses the memory references from that profile to prefetch the memory objects associated with state 4. In the example of FIG. 12, the predictor 46 then determines that, again based on past performance, the next state (i.e., state 5) is also 100% deterministic. Accordingly, the prefetcher 304 accesses the memory profile of state 5 and uses the memory references from that profile to prefetch the memory objects associated with state 5.

**[0083]** In the example of FIG. 12, state 5 marks the end of the macroscopic transaction and, thus, has a high entropy value. Accordingly, the next program state may not be predicted with a high degree of certainty. As explained above, depending on the prefetching strategy selected, the prefetcher 304 may respond to the occurrence of a state with a high entropy value in any number of ways. For example, it may prefetch some, none or all of the next possible states.

**[0084]** Irrespective of the prefetching strategy chosen, program execution continues. In the example of FIG. 12 program execution proceeds from state 5 to state 18, which marks the start of a new macroscopic transaction. Upon entering state 18, the predictor 46 and the prefetcher 304 operate as explained above to predict the next most probable state(s) (e.g., state 19) and to prefetch some, all or none of the memory objects associated with those state(s) depending on the prefetching strategy and, possibly, the entropy value of the current state.

**[0085]** Flowcharts representative of example machine readable instructions for implementing the apparatus 300 of FIG. 10 are shown in FIGS. 8, 13A-13C, 14 and/or 15. In this example, the machine readable instructions comprise a program for execution by a processor such as the processor 1012 shown in the example computer 1000 discussed below in connection with FIG. 16. The program may be embodied in software stored on a tangible medium such as a CD-ROM, a floppy disk, a hard drive, a digital versatile disk (DVD), or a memory associated with the processor 1012, but persons of ordinary skill in the art will readily appreciate that the entire program and/or parts thereof could alternatively be executed by a device other than the processor 1012 and/or embodied in firmware or dedicated hardware in a well known manner. For example, any or all of the trace sampler 12, the state identifier 14, the predictor 46, the memory state monitor 302 and/or the prefetcher 304 could be implemented by software, hardware, and/or firmware. Further, although the example program is described with reference to the flowcharts illustrated in FIGS. 8, 13A-13C, 14 and/or 15, persons of ordinary

skill in the art will readily appreciate that many other methods of implementing the example apparatus 300 may alternatively be used. For example, the order of execution of the blocks may be changed, and/or some of the blocks described may be changed, eliminated, or combined.

**[0086]** Since, as explained above, some of the structures of the example apparatus 10 are substantially identical to structures of the example apparatus 300, if those overlapping structures are implemented via software and/or firmware, they may be implemented by similar programs. Thus, for example, the trace sampler 12, the state identifier 14 and the predictor 46 may be implemented in the example apparatus 300 using substantially the same machine readable instructions described above in connection with FIGS. 8 and 9A-9C. In the interest of brevity, the blocks of the program used to implement the apparatus 300 which are the same or substantially the same as the blocks of the program used to implement the apparatus 10, will be described in abbreviated form here. The interested reader is referred to the above description for a full description of those blocks. To facilitate this process, like blocks are labeled with like reference numerals in FIGS 9A-9C and 13A-13C.

**[0087]** As mentioned above, the trace sampler 12 is implemented by substantially the same machine readable instructions in the apparatus 300 as in the apparatus 10. Thus, the above-description of blocks 100-108 applies to both the example apparatus 10 and the apparatus 300, except that in the apparatus 300, the trace sampler 12 generates a memory address trace. The program of FIG. 8 begins at block 100 where the target program begins

execution. While the target program executes, the trace sampler 12 creates a memory address trace (block 102). Control proceeds from block 102 to block 104.

**[0088]** If a trace processing thread has already been invoked (block 104), control proceeds from block 104 to block 106. If the trace 18 of the program is complete (block 106), the program of FIG. 8 terminates. Otherwise, if the trace 18 of the program is not complete (block 106), control returns to block 102 where the recording of the trace 18 continues.

**[0089]** If the trace processing thread has not already been invoked (block 104), control proceeds to block 108. At block 108, the trace processing thread is initiated. Control then proceeds to block 106. Control continues to loop through blocks 100-108 until the target program stops executing and the trace 18 is complete.

**[0090]** Once a trace processing thread is spawned (block 108, FIG. 8), the illustrated trace processing thread begins at block 120 (FIG. 13A) where the signature developer 16 obtains a set 26 of entries from the trace 18 created by the trace sampler 12. Once the entries to create a set 26 are retrieved from the trace 18 (block 120), the weight assigning engine 34 adjusts the values of the retrieved entries such that later entries are given greater weight than earlier entries (block 122). Once the values of the entries have been weighted by the weight assigning engine 34 (block 122), the signature developer 16 maps the entries in the set 26 to an n-bit vector to create a possible state signature 28 for the set 26 (block 124). After the possible state signature 28 is generated (block 124), the state distinguisher 38 determines whether the possible state signature

28 is the first possible state signature (block 126). If it is the first possible state signature (block 126), the first possible state signature is, by default, defined to be the first state signature. Thus, the state distinguisher 38 sets a current state signature variable equal to the possible state signature 28 (block 128) and creates a state data structure in the state array 44 for the first state (block 130). The memory state monitor 302 then writes the memory references of the memory profile associated with the current state in the state data structure (block 331).

**[0091]** The signature developer 16 then collects the next set 26 of entries for creation of a possible state signature 28 (block 132). In the illustrated example, the sets 26 used by the signature developer 16 to create the possible signatures 28 are overlapping. Thus, the signature developer 16 may create the next set 26 of entries by dropping the oldest entr(ies) from the last set 26 of entries and adding a like number of new entr(ies) to create a new current set 26 (block 132). Control then returns to block 122 (FIG. 13A) where the entries in the new current set are weighted as explained above.

**[0092]** When at block 126 of FIG. 13A, the current possible state signature is not the first possible state signature, control will skip from block 126 to block 134 (FIG. 13B). At block 134, the state distinguisher 38 calculates the difference between the current state signature, and the current possible state signature. The state distinguisher 38 then compares the computed difference to a threshold. If the computed difference exceeds the threshold (block 136), a program state change has occurred and control proceeds to block 138. If the computed difference does not exceed the

threshold (block 136), the signature developer 16 collects the next set 26 of entries for creation of a possible state signature 28 (block 132, FIG. 13A) and control returns to block 122 as explained above. By adjusting the threshold, one may adjust the duration and number of the program states. It is, thus, this threshold that may be adjusted to ensure the prefetching operation is performed at an appropriate time as explained above.

**[0093]** Assuming for purposes of discussion that a program state change has occurred (block 136 of FIG. 13B), the state distinguisher 38 sets the current state signature variable equal to the current possible state signature 28 (block 138). The state distinguisher 38 then examines the signatures present in the state array 44 to determine if the current state signature corresponds to the signature of a known state (block 140). If the current state signature is a known state signature, control advances to block 160 (FIG. 13C). Otherwise, if the current state signature is not a known state signature (i.e., the current state signature does not correspond to a state already existing in the state array 44), control advances to block 142 (FIG. 13B).

**[0094]** Assuming for purposes of discussion that the current state signature is not a known state signature (e.g., the current program state is a new program state) (block 140), the state distinguisher 38 creates a state data structure in the state array 44 for the first state (block 142) as explained above in connection with block 130. The memory state monitor 302 then writes the memory references of the memory profile associated with the current state in the state data structure (block 343).



**[0095]** The state transition monitor 48 then updates the last state's probability sub-array to reflect the transition from the last state to the new current state (block 144). Control then proceeds to block 146 where the state distinguisher 38 determines if the state array 44 has become full. If the state array 44 is not full (block 146), control returns to block 132 of FIG. 13A. If the state array is full (block 146), control advances to block 150 (FIG. 13B) where the state distinguisher 38 deletes the stalest state data structure from the state array 44. Once the stalest state data structure is eliminated (block 150), control returns to block 132 of FIG. 13A.

**[0096]** Assuming that the current state signature is a known state signature (block 140 of FIG. 13B), control proceeds to block 358 (FIG. 13C). At block 358, the memory state monitor 302 updates the memory profile of the current state. For example, the memory state monitor 302 may filter the memory profile by adding, deleting, and/or changing one or more of the memory references in the memory profile to reflect the memory references most recently associated with the program state as explained above.

**[0097]** The state transition monitor 48 then updates the last state's probability sub-array to reflect the transition from the last state to the new current state (block 160). Control then proceeds to block 162 where the entropy calculator 50 calculates the entropy value of the current state.

**[0098]** Once the entropy value is calculated (block 162), the event predictor 54 identifies the next most probable state(s) (block 164). The prefetcher 304 then executes the prefetching strategy of the apparatus 300. As explained above, there are many different prefetching strategies that may be

employed by the prefetcher 304. For example, as shown in FIG. 14, the prefetcher 304 may always prefetch the memory objects for the next most probable state or a set of the next most probable states identified by the predictor 46 (block 380). Alternatively, as shown in FIG. 15, the prefetching activity of the prefetcher 304 may be dependent upon the entropy value calculated for the current program state.

**[0099]** In the example of FIG. 15, the prefetcher 304 first retrieves and compares the entropy value of the current program state to a threshold X (block 382). If the entropy value of the current state is below the threshold X (block 382), the prefetcher 304 prefetches the memory objects associated with the most probable state or states (again, depending on the strategy employed) (block 384). If, however, the entropy value of the current state is above the threshold X (block 382), the prefetcher 304 compares the entropy value of the current program state to a threshold Y (block 386). If the entropy value of the current state is below the threshold Y (block 386), the prefetcher 304 prefetches the memory objects associated with all of the known next probable states (block 388). If, however, the entropy value of the current state is above the threshold Y (block 386), the prefetcher 304 does not prefetch any memory objects at this time.

**[00100]** Irrespective of the prefetching strategy employed, after the prefetching strategy is executed (block 366), control advances to block 168 (FIG. 13C). At block 168, the event predictor 54 examines the entropy values of the last few states to determine if an entropy spike has occurred (block 168). If an entropy spike is identified (block 168), the event predictor 54

identifies the program state corresponding to the entropy spike as the last state of a macroscopic transaction or sub-transaction (block 170). If an entropy spike is not identified (block 168), the end of a macroscopic transaction or sub-transaction has not occurred. Accordingly, control skips block 170 and returns to block 132 (FIG. 13A).

**[00101]** FIG. 16 is a block diagram of an example computer 1000 capable of implementing the apparatus and methods disclosed herein. The computer 1000 can be, for example, a server, a personal computer, a personal digital assistant (PDA), an Internet appliance, a DVD player, a CD player, a digital video recorder, a personal video recorder, a set top box, or any other type of computing device.

**[00102]** The system 1000 of the instant example includes a processor 1012. For example, the processor 1012 can be implemented by one or more Intel® microprocessors from the Pentium® family, the Itanium® family, the XScale® family, or the Centrino™ family. Of course, other processors from other families are also appropriate.

**[00103]** The processor 1012 is in communication with a main memory 306 including a volatile memory 1014 and a non-volatile memory 1016 via a bus 1018. The volatile memory 1014 may be implemented by Synchronous Dynamic Random Access Memory (SDRAM), Dynamic Random Access Memory (DRAM), RAMBUS Dynamic Random Access Memory (RDRAM) and/or any other type of random access memory device. The non-volatile memory 1016 may be implemented by flash memory and/or any other desired type of memory device. Access to the main memory 1014,

1016 is typically controlled by a memory controller (not shown) in a conventional manner.

**[00104]** The computer 1000 also includes a conventional interface circuit 1020. The interface circuit 1020 may be implemented by any type of well known interface standard, such as an Ethernet interface, a universal serial bus (USB), and/or a third generation input/output (3GIO) interface.

**[00105]** One or more input devices 1022 are connected to the interface circuit 1020. The input device(s) 1022 permit a user to enter data and commands into the processor 1012. The input device(s) can be implemented by, for example, a keyboard, a mouse, a touch screen, a track-pad, a trackball, isopoint and/or a voice recognition system.

**[00106]** One or more output devices 1024 are also connected to the interface circuit 1020. The output devices 1024 can be implemented, for example, by display devices (e.g., a liquid crystal display, a cathode ray tube display (CRT), a printer and/or speakers). The interface circuit 1020, thus, typically includes a graphics driver card.

**[00107]** The interface circuit 1020 also includes a communication device such as a modem or network interface card to facilitate exchange of data with external computers via a network 1026 (e.g., an Ethernet connection, a digital subscriber line (DSL), a telephone line, coaxial cable, a cellular telephone system, etc.).

**[00108]** The computer 1000 also includes one or more mass storage devices 1028 for storing software and data. Examples of such mass

storage devices 1028 include floppy disk drives, hard drive disks, compact disk drives and digital versatile disk (DVD) drives. The mass storage device 1028 may implement the memory 44.

**[00109]** As an alternative to implementing the methods and/or apparatus described herein in a system such as the device of FIG. 16, the methods and/or apparatus described herein may alternatively be embedded in a structure such as processor and/or an ASIC (application specific integrated circuit).

**[00110]** From the foregoing, persons of ordinary skill in the art will appreciate that the above disclosed methods and apparatus may be implemented in a static compiler, a managed run-time environment just-in-time compiler (JIT), and/or directly in the hardware of a microprocessor to achieve performance optimization in executing various programs and/or in memory operations associated with an executing program. In the context of the apparatus 300, a static compiler could exploit the predictable repetitive behavior by generating speculative threads to prefetch the memory objects associated with the next probable program states. Similarly, an MRTE JIT (managed run time environment just in time) engine could use the above disclosed methodology to prefetch memory objects based on dynamic profiling.

**[00111]** Although certain example methods, apparatus, and articles of manufacture have been described herein, the scope of coverage of this patent is not limited thereto. On the contrary, this patent covers all

methods, apparatus and articles of manufacture fairly falling within the scope of the appended claims either literally or under the doctrine of equivalents.